

Hexadecapole Approximation in Planetary Microlensing

Andrew Gould

Department of Astronomy, Ohio State University, 140 W. 18th Ave., Columbus, OH 43210, USA; gould@astronomy.ohio-state.edu

ABSTRACT

The frequency of microlensing planet detections, particularly in difficult-to-model high-magnification events, is increasing. Their analysis can require tens of thousands of processor hours or more, primarily because of the high density and high precision of measurements whose modeling requires time-consuming finite-source calculations. I show that a large fraction of these measurements, those that lie at least one source diameter from a caustic or the extension from a cusp, can be modeled using a very simple hexadecapole approximation, which is one to several orders of magnitude faster than full-fledged finite-source calculations. Moreover, by restricting the regions that actually require finite-source calculations to a few isolated “caustic features”, the hexadecapole approximation will, for the first time, permit the powerful “magnification-map” approach to be applied to events for which the planet’s orbital motion is important.

Subject headings: gravitational lensing – planetary systems – methods numerical

1. Introduction

Microlensing planets are being discovered at an accelerating rate, with one being reported in 2003 (Bond et al. 2004), three in 2005 (Udalski et al. 2005; Beaulieu et al. 2006; Gould et al. 2006), two in 2006 (Gaudi et al. 2008), and perhaps as many as six in 2007. It has been a huge challenge for modelers to keep up with these discoveries, in large part because the computing requirements are often daunting: the parameter space is large, the χ^2 surface is complex and generally contains multiple minima, and the magnification calculations are computationally intensive. Proper treatment of individual events can require tens of thousands processor hours, or more. Indeed, some potential planetary events have still not been fully modeled because of computational challenges.

Planetary microlensing events are recognized through two broad channels, one in which the lightcurve perturbation is generated by the so-called “planetary caustic” that is directly

associated with the planet (Gould & Loeb 1992) and the other in which it is generated by the “central caustic” that is associated with the host star (Griest & Safizadeh 1998). The former events are relatively easy to analyze and indeed the event parameters can be estimated reasonably well by inspection of the lightcurve. The latter events are generally much more difficult.

There are several interrelated reasons for this. First, planets anywhere in the system can perturb the central caustic. This is why these events are avidly monitored, but by the same token it is often not obvious without an exhaustive search which planetary geometry or geometries are responsible for the perturbation. Second, this very fact implies that several members of a multi-planet system can be detected in central-caustic events (Gaudi et al. 1998). Multiple planets create a larger, more complicated parameter space, which can increase the computation time by a large factor. Even if there are no obvious perturbations caused by a second planet, an exhaustive search should be conducted to at least place upper limits on their presence. Third, if the source probes the central caustic, it is ipso facto highly magnified. Such events are brighter and more intensively monitored than typical events and so have more and higher-quality data. While such excellent data are of course a boon to planet searches, they also require more and more-accurate computations, which requires more computing time. Fourth, central-caustic planetary events are basically detectable in proportion to the size of the caustic, which roughly scales $\propto q/|b - 1|$, where q is the planet/star mass ratio and b is the planet-star separation in units of the Einstein ring. Thus, these events are heavily biased toward planets with characteristics that make the caustic big. Such big caustics can undergo subtle changes as the planet orbits its host, and in principle these effects can be measured, thus constraining the planet’s properties. Exploration of these subtle variations requires substantial additional computing time. Finally, there is also a bias toward long events, simply because these unfold more slowly and so increase the chance that they will be recognized in time to monitor them intensively over the peak. Such long events often display lightcurve distortions in their wings due to the Earth’s orbital motion which, if measured, can further constrain the planet properties. However, probing this effect (called “microlens parallax”) requires yet another expansion of parameter space. Moreover, the “parallax” signal must be distinguished from “xallarap” (effects of the source orbiting a companion), whose description requires a yet larger expansion of parameter space.

There are two broad classes of binary-lens (or triple-lens) magnification calculations: point-source and finite-source. The former can be used whenever the magnification is essentially constant (or more precisely, well-characterized by a linear gradient) over face of the source, while the latter must be used when this condition fails. Point-source magnifications can be derived from the solution of a 5th (or 10th) order complex polynomial equation and

are computationally very fast. When the point-source is well separated from the caustics (as it must be to satisfy the linear-gradient condition) then this calculation is also extremely robust and accurate.

The main computational challenge in modeling planetary events comes from the finite-source calculations. Almost all integration schemes use inverse-ray shooting, which avoids all the pathologies of the caustics by performing an integration over the *image* plane (where the rays behave smoothly) and simply asks which of the rays land on the source. The problem is that a large number of rays must be “shot” to obtain an accurate estimate of the magnification, which implies that high-quality data demand proportionately longer computations. Of course, the higher the magnification, the larger the images, and so the more rays are required. There are various schemes to expedite inverse ray-shooting, including clever algorithms for identifying the regions that must be “shot” and pixelation of the source plane. The bottom line is, however, that the overwhelming majority of computation time is spent on finite-source calculations.

Here I present a third class of binary-lens (or triple-lens) computation that is intermediate between the two classes just described: the hexadecapole approximation. Pejcha & Heyrovský (2007) expand finite-source magnification to hexadecapole order and illustrate that the quadrupole term by itself can give quite satisfactory numerical results. In this paper, I develop a simple prescription for evaluating this expansion. While this algorithm is about 10 times slower than point-source calculations, it is one to several orders of magnitude faster than finite-source calculations. The method can be applied whenever the source center is at least two source radii from a caustic or the extension from a cusp. Typically, well over half the non-point-source points satisfy this condition, meaning that the method can reduce computation times by a factor of several. Moreover, by isolating the small regions of the lightcurve where finite-source calculations must be used, the method opens up the possibility that the finite-source computations themselves can be radically expedited for the special, but very interesting, class of events in which planetary orbital motion is measured.

2. Hexadecapole Approximation

If the source does not straddle a caustic, then the magnification field can be Taylor expanded around the source center (x_0, y_0) as a function of coordinate position (x, y) (all distances being expressed in units of the Einstein radius),

$$A(x, y) = \sum_{n=0}^{\infty} \sum_{i=0}^n A_{ni} (x - x_0)^i (y - y_0)^{n-i}. \quad (1)$$

We wish to evaluate $A_{\text{finite}}(\rho)$, the magnification of a source of radius ρ ,

$$A_{\text{finite}}(\rho, x_0, y_0) \equiv \frac{\int_0^\rho dw \int_0^{2\pi} d\eta A(w, \eta)}{\pi \rho^2}, \quad (2)$$

in terms of the A_{ni} . Here (w, η) are polar coordinates: $(w \cos \eta, w \sin \eta) \equiv (x - x_0, y - y_0)$. To do so, we first average $A(x, y)$ over a ring of radius w and obtain,

$$A(w) = \frac{\int_0^{2\pi} d\eta A(w \cos \eta, w \sin \eta)}{2\pi} = A_0 + A_2 w^2 + A_4 w^4 + \dots \quad (3)$$

where

$$A_0 \equiv A_{00}; \quad A_2 \equiv \frac{A_{20} + A_{22}}{2}; \quad A_4 \equiv \frac{3A_{40} + A_{42} + 3A_{44}}{8}. \quad (4)$$

To evaluate A_{finite} in terms of $A_0, A_2, A_4 \dots$, we must first specify a limb-darkening law for the surface brightness $S(w)$, which for simplicity we take to be linear,

$$S(w) = \left[1 - \Gamma \left(1 - \frac{3}{2} \sqrt{1 - \frac{w^2}{\rho^2}} \right) \right] \frac{F}{\pi \rho^2}, \quad (5)$$

where Γ is the limb-darkening coefficient and F is the source flux. We then substitute into equation (3) to obtain

$$A_{\text{finite}} \equiv \frac{\int_0^\rho dw 2\pi w A(w) S(w)}{F} = A_0 + \frac{A_2 \rho^2}{2} \left(1 - \frac{1}{5} \Gamma \right) + \frac{A_4 \rho^4}{3} \left(1 - \frac{11}{35} \Gamma \right) + \dots \quad (6)$$

Note that the “limb-darkening factors” in parentheses are simply $(2\pi/F) \int_0^\rho dw S(w) w^{2n+1}$ (Pejcha & Heyrovský 2007), and therefore do not depend on the magnification in any way. Hence, for any given adopted limb-darkening profile, these can be calculated just once.

Let us now assume that the field is adequately described by a hexadecapole. Averaging over the four points on a w -ring that are shifted by an arbitrary angle ϕ relative the cardinal directions, we obtain

$$\begin{aligned} A_{w,+} &\equiv \frac{1}{4} \sum_{j=0}^3 A \left[w \cos \left(\phi + j \frac{\pi}{2} \right), w \sin \left(\phi + j \frac{\pi}{2} \right) \right] - A_0 \\ &= A_2 w^2 + \frac{(A_{40} + A_{44})(1 + \cos^2 2\phi) + A_{42} \sin^2 2\phi}{4} w^4, \end{aligned} \quad (7)$$

while rotating this geometry by $\pi/4$ gives,

$$A_{w,\times} = A_2 w^2 + \frac{(A_{40} + A_{44})(1 + \sin^2 2\phi) + A_{42} \cos^2 2\phi}{4} w^4. \quad (8)$$

For a given source of size ρ and position (x_0, y_0) , one can therefore determine A_0 , $A_{\rho,+}$, $A_{\rho/2,+}$, and $A_{\rho,\times}$, from a total of 13 point-source calculations, and thus derive

$$A_2\rho^2 = \frac{16A_{\rho/2,+} - A_{\rho,+}}{3}, \quad A_4\rho^4 = \frac{A_{\rho,+} + A_{\rho,\times}}{2} - A_2\rho^2, \quad (9)$$

which can then be substituted into equation (6) to obtain $A_{\text{finite}}(\rho, x_0, y_0)$.

3. Range of Validity

Clearly, this approximation cannot be used when the source lies on a caustic, but how close can the source be before the approximation breaks down? The breakdown will be driven by the leading term of the caustic’s singularity, so it is sufficient to examine idealized cases whose leading-order behavior is the same as that of real caustics. More formally, one can write the magnification field as the linear sum of an idealized singularity and a more complicated, but well-behaved field. Only the former will contribute to the breakdown of the approximation. Since equations (6), (7), (8), and (9) are strictly linear, this decomposition is absolutely rigorous. Caustic singularities come in two basic varieties, fold caustics and cusps. The former diverge as $(\Delta u_{\perp})^{-1/2}$ as one approaches the caustic from the inside, where Δu_{\perp} is the perpendicular distance to the fold. The latter are much more complex. They diverge as $(\Delta u)^{-1}$ in the immediate neighborhood of the cusp as one approaches it from the outside. However, at greater distances, they develop into long “fingers” (Gould & Loeb 1992; Pejcha & Heyrovský 2007).

3.1. Fold Caustic

For simplicity, I consider a uniform source that lies a distance $z\rho$ (where ρ is the source radius) from the fold. I begin by assuming that the magnification is dominated by the two “new images” that meet on the critical curve. I then discuss how the result is changed when this assumption is relaxed. The magnified flux is then given by

$$\frac{A(z)}{A_0} = \frac{2}{\pi} \int_{-1}^1 dx \sqrt{\frac{1-x^2}{1+x/z}} = \frac{2}{\pi} \int_{-1}^1 dx \sqrt{1-x^2} \sum_{m=0}^{\infty} \frac{(2m-1)!!}{2^m m!} \left(\frac{x}{z}\right)^m, \quad (10)$$

which after some algebra [e.g., $\int_0^1 dy y^m (1-y)^n = m!n!/(m+n+1)!]$, simplifies to

$$\frac{A(z)}{A_0} = \sum_{n=0}^{\infty} \frac{(4n-1)!!(2n-1)!!}{2^{3n}(2n)!(n+1)!} z^{-2n}, \quad (11)$$

where $n = m/2$. That is,

$$\frac{A(z)}{A_0} = 1 + \frac{3}{2^5} z^{-2} + \frac{35}{2^{10}} z^{-4} + \frac{1155}{2^{16}} z^{-6} + \frac{45045}{2^{22}} z^{-8} + \dots \quad (12)$$

Hence the error due to the hexadecapole approximation is

$$\frac{\Delta A(z)}{A_0} = \frac{1155}{2^{16}} z^{-6} = 2.8 \times 10^{-4} \left(\frac{z}{2}\right)^{-6}. \quad (13)$$

If one is “sufficiently near” a fold caustic, the magnification will always be dominated by the two “new images”. However, for planetary caustics, this will generally no longer be the case even at one or two source radii from the caustic. The net effect is to “dilute” the caustic divergence and so to make the hexadecapole approximation valid at even closer separations than indicated by equation (13). See § 5.

3.2. Cusp

For simplicity, I analyze the case of a $(\Delta u)^{-1}$ divergence and then discuss how the results may be expected to change for real cusps. In this case, we have

$$\begin{aligned} \frac{A(z)}{A_0} &= \frac{1}{2\pi} \int_0^1 dx^2 \int_0^{2\pi} d\theta \left(1 + \frac{2x \cos \theta}{z} + \frac{x^2}{z^2}\right)^{-1/2} \\ &= \frac{1}{2\pi} \int_0^1 dx^2 \int_0^{2\pi} d\theta \sum_{m=0}^{\infty} (-1)^m \frac{(2m-1)!!}{m! 2^m} \sum_{i=0}^m C_i^m \left(\frac{2x \cos \theta}{z}\right)^i \left(\frac{x^2}{z^2}\right)^{m-i}, \end{aligned} \quad (14)$$

which after some algebra [e.g., $\langle (2 \cos \theta)^{2k} \rangle = C_k^{2k}$] reduces to

$$\frac{A(z)}{A_0} = \sum_{n=0}^{\infty} \frac{z^{-2n}}{(n+1)} \sum_{j=0}^n \left(-\frac{1}{2}\right)^{j+n} \frac{(2n+2j-1)!!}{(n-j)! [j!]^2}, \quad (15)$$

where, again, $n = m/2$. That is,

$$\frac{A(z)}{A_0} = 1 + \frac{1}{2^3} z^{-2} + \frac{3}{2^6} z^{-4} + \frac{25}{2^{10}} z^{-6} + \frac{245}{2^{14}} z^{-8} + \dots, \quad (16)$$

so that the error due to the hexadecapole approximation is,

$$\frac{\Delta A(z)}{A_0} = \frac{25}{2^{10}} z^{-6} = 3.8 \times 10^{-4} \left(\frac{z}{2}\right)^{-6}. \quad (17)$$

Equation (17) gives a reasonable lower limit on how closely one may approach a cusp using the hexadecapole approximation. However, because the cusp develops a linear, finger-like structure at moderate distances, this approximation can fail well away from the cusp. Nevertheless, I find numerically that the total duration of the failure is generally represented reasonably well by equation (17). See § 5.

Thus, for both fold caustics and cusps, the hexadecapole approximation can reduce the error below 0.1%, except for intervals characterized by $z = 2.5$.

4. Implementation

The main consideration when implementing this approximation is to make certain that it is applied only in its range of validity. This is easiest when one is probing an already-located minimum, which is often the most time-consuming part of the investigation. One can then simply plot the difference between the hexadecapole approximation and the finite-source calculation for a *single model* as a function of time, and so locate empirically the regions of the former’s range of validity. If a “safety zone” is placed around the finite-source regions, then there is little danger that it will be crossed during the minor excursions that occur while probing a minimum. It is straightforward to determine automatically whether such unexpected crossings are occurring simply by comparing the finite-source and hexadecapole calculations for the first and last points of each finite-source region. If these do not agree, the “safety zone” has been crossed.

One must be more careful when applying this method to blind searches because it is harder to determine whether any particular stretch of the lightcurve is either crossing or nearby a caustic. In some cases, this will be straightforward and in others more difficult. The one general point to note is that the same “safety zone” check can be made.

Finally, I remark that in some of the lightcurve regions where the hexadecapole approximation is valid, it may be overkill: the quadrupole or even monopole (point-source) approximations may be perfectly satisfactory. Again, it is straightforward to find these subregions, simply by mapping the hexadecapole/quadrupole and hexadecapole/monopole differences for a single model.

5. An Illustration of the Method

Figures 1 and 2 give a practical example of the hexadecapole approximation. The top panel of Figure 1 shows a caustic geometry (*black*) that is based loosely on the caustic in the planetary event OGLE-2005-BLG-071 (Udalski et al. 2005), but with a different source trajectory (*blue*) that has been chosen to maximize the number of illustrative “features”: the source passes by two cusps and then enters and exits the caustic.

The middle panel shows the resulting lightcurve (*black*) together with three successive levels of approximation: monopole (i.e., point-source *blue*), quadrupole (*red*), and hexadecapole (*green*). The bottom panel shows the residuals of the latter three relative to the first.

There are several points to note. First, the hexadecapole approximation works extremely well over the entire region shown except for a few source-radius crossing times in the immediate vicinity of the first cusp approach and the two caustic crossings. Second, the

point-source approximation basically does not work at all over the entire region shown, at least if one is attempting to achieve precisions of 0.1% (which is typically required). Third, there are significant regions where the quadrupole approximation is adequate.

Figure 2 is a zoom of Figure 1, focusing on the region of the two caustic crossings. The crosses in the bottom panel show the predictions for the breakdown of the quadrupole (*red*) and hexadecapole (*green*) approximation inside the caustic (eqs. [12] and [13]), assuming one is trying to achieve a precision of 0.05%. The *black* crosses indicate the prediction for breakdown outside the caustic, namely one source radius from the caustic.

Both the quadrupole and hexadecapole approximations prove to be too conservative in that the approximations remain valid substantially closer to the caustic than expected. The reason for this is clear from the middle panel: at the points of the predicted breakdown, the underlying magnification profile is no longer dominated by the square-root singularity that is produced by the two “new images”. Recall from § 3.1 that equations (12) and (13) were specifically derived under the assumption these images would dominate the magnification.

Returning to Figure 1, the situation is more complex for the cusp approaches. Equation (17) predicts that the hexadecapole approximation will break down from -25.8 to -22.7 . The actual breakdown is from -27.6 to -23.2 . This displacement to the left reflects the “finger” of enhanced magnification that extends from the cusp axis, as indicated in the top panel. In fact, I find that for trajectories passing farther from this caustic, the breakdown occurs at even earlier times, again corresponding to the intersection of the trajectory with the cusp axis (rather than the point of closest approach). Moreover, the breakdown continues to occur even when the point of closest approach is well beyond the separation predicted by equation (17). Nevertheless, the full duration of the breakdown is never much greater than $2z$ as predicted by this equation. In brief, for cusps, the hexadecapole approximation does work except for brief intervals whose duration is given by equation (17), but the time of breakdown is not accurately predicted by this equation. Similar remarks apply to the quadrupole approximation.

6. Application to Lenses With Orbital Motion

While the hexadecapole approximation can save substantial computation time in a wide range of cases, it may be especially useful for planetary lenses with measurable orbital motion by rendering them accessible to the “magnification map” technique (Dong et al. 2006). Magnification maps are potentially very powerful. One first chooses “map parameters”, i.e., (b, q) for a single planet or $(b_1, q_1, b_2, q_2, \phi)$ for a triple system, where ϕ is the angle between the two planets. Then one shoots the entire Einstein ring (out to a specified width corresponding to, say, magnification $A = 100$). One then both stores the individual rays and

tiles the source plane with hexagonal pixels, keeping track of the number of rays landing in each pixel. Pixels landing wholly within the source are evaluated at their centroid, while pixels that cross the boundary are evaluated on a ray-by-ray basis. Using this map, one can minimize over the remaining microlensing variables (time of closest approach t_0 , impact parameter u_0 , Einstein timescale t_E , source trajectory angle α , source size ρ , as well as parallax and xallarap parameters, if these are needed). Each map can be created in a few seconds and fully explored in a few minutes, thereby permitting a rapid Markov-chain walk through map space.

The drawback is that, to date, magnification maps have not been applicable to lenses with significant orbital motion: to the extent that the lens separation changes during the event, different maps would be needed at different phases of the event, potentially a very large number of them. In some cases, this problem is now completely resolved. For example, OGLE-2005-BLG-071 exhibits some signature of rotation (Dong et al. 2008), which had been difficult to probe simultaneously with finite-source effects. Because the source trajectory comes no closer than $z = 10$ source radii from the cusp, this event can now be handled completely in the hexadecapole approximation.

However, even for events with one or several caustic features that require finite-source calculations, it will now be possible to evaluate these with maps. Each caustic feature lasts only about the time required for the source to cross its own diameter, which is typically a few hours. The orbital motion during these features is negligible, implying that the lens geometry can be adequately represented by a single map. Several maps can be created, one for each feature occurring at different times. The points “between features” can be evaluated in the hexadecapole approximation, which allows a continuously evolving lens geometry.

7. Conclusions

I have identified an intermediate regime between the ones where finite-source effects are dominant and negligible, respectively. In this regime, magnifications can be evaluated with very high precision using a simple hexadecapole approximation, for which I give a specific prescription. Outside of small (few source-diameter crossing times) regions associated with caustic crossings and cusp approaches, the approximation has a fractional error of well under 0.1%. Some events can now be analyzed without any traditional finite-source calculations, while for others these calculations will be drastically reduced. In particular, by restricting the regions that do absolutely require finite-source calculation to a few isolated zones, this approximation opens the possibility of applying “magnification maps” to planetary systems experiencing significant orbital motion.

I thank Subo Dong for seminal discussions. The paper benefited greatly from the suggestions of an anonymous referee. This work was supported by NSF grant AST 042758.

REFERENCES

- Beaulieu, J.-P. et al. 2005, *Nature*, 439, 437
- Bond, I.A., et al. 2004, *ApJ*, 606, L155
- Dong, S., et al. 2006, *ApJ*, 642, 842
- Dong, S., et al. 2008, in preparation
- Gaudi, B. S., Naber, R.M. & Sackett, P. D. 1998, *ApJ*, 502, L33
- Gaudi, B.S., et al. 2008, *Science*, 315, 927
- Gould, A., & Loeb, A. 1992, *ApJ*, 396, 104
- Gould, A., et al. 2004, *ApJ*, 644, L37
- Griest, K. & Safizadeh, N. 1998, *ApJ*, 500, 37
- Pejcha, O. & Heyrovský, D. 2007, *ApJ*, submitted (astro-ph/0712.2217)
- Udalski, A., et al. 2005, *ApJ*, 628, L109




Fig. 1.— Top: source trajectory (*blue*) through caustic geometry (*black*) of a simulated high-magnification microlensing event, in units of the source radius. Middle: resulting lightcurve as found from full finite-source calculation (*black*) and the monopole (*blue*), quadrupole (*red*), and hexadecapole (*green*) approximations. Bottom: residuals of the three approximations relative to the full calculation. See § 5 for full discussion.




Fig. 2.— Zoom of Fig. 1, focusing on region of the two caustic crossings. Crosses in middle panel show predictions for range of validity of the quadrupole (*red*) and hexadecapole (*green*) approximations inside the caustic and for both (*black*) outside. See § 5 for full discussion.